IN THE UNITED STATES PATENT AND TRADEMARK OFFICE BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of:		§	
	Curtis R. PRIEM	§	Confirmation No.: 9182
		§	
Serial No.:	10/804,945	§	Group Art Unit: 2111
		§	
Filed:	March 19, 2004	§	Examiner: Thomas J. Cleary
		§	
For:	METHOD AND APPARATUS FOR LATENCY BASED THREAD SCHEDULING	§	

MAIL STOP APPEAL BRIEF-PATENTS
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

REPLY BRIEF

Dear Sir:

The Appellants submit this Reply Brief to the Board of Patent Appeals and Interferences in response to the Examiner's Answer mailed on January 13, 2009. While Applicants maintain each of the arguments submitted in Applicants' previously submitted Appeal Brief, Applicants make the following further arguments in light of the Examiner's Answer. Although Applicants believe that no additional fees are due in connection with this reply, the Commissioner is hereby authorized to charge Deposit Account No. 20-0782/NVDA/P000455/SW for any fees necessary to make this reply timely and acceptable to the Office.

REMARKS

Regarding claims 1 and 9:

Independent claims 1 and 9 recite the limitations of (i) masking interrupts from hardware devices, (ii) unmasking interrupts from the hardware devices, (iii) simultaneously rearranging threads in a single queue, and (iv) the order in which the threads will be serviced in the single queue is simultaneously rearranged to schedule the thread for processing. The single queue represents the order in which all threads will be serviced for all of the hardware devices. Zolnowsky, Jones, Browning, and Ramakrishnan fail to teach or suggest these limitations.

Modifying the architecture of Zolnowsky to rearrange the order in which the threads will be serviced in a single queue, conflicts with the stated qualities of Zolnowsky's architecture. It would not have been obvious to a person of ordinary skill in the art to modify the scheduling mechanism of Zolnowsky to incorporate the use a single queue for all requests from all hardware devices. Zolnowsky teaches a multiple processor system where requests are distributed among many queues. The mechanism taught by Zolnowsky relies on providing a dispatch queue for each processor to minimize lock contention between the multiple processors and provide the desired processing throughput (see col. 5, lines 42-44 and col. 2, lines 3-9). In addition to the per-processor dispatch queues, Zolnowsky teaches using a global dispatch queue for higher priority real time threads to ensure that the real time threads are dispatched as quickly as possible (see col. 5, lines 45-54). Thus, modifying the architecture of Zolnowsky to rearrange the order in which the threads will be serviced in a single queue, is in direct conflict with Zolnowsky's use of a separate global dispatch queue for higher priority real time threads and a dispatch queue for each processor to minimize lock contention. In short, Zolnowsky explicitly teaches away from the modification that the Examiner is suggesting and is not an appropriate 103 (a) reference.

Additionally, none of the cited references teaches or suggests that the limitations of ordering all requests from all of the hardware devices in a single queue and that all of the threads are simultaneously rearranged in the single queue. Instead, the cited references describe various mechanisms that use multiple queues and selecting threads from the queues to schedule the threads for processing.

Zolnowsky relies on selecting threads from the queue rather than simultaneously rearranging the order in which the threads will be serviced in a single queue. As recited in claim 2 of Zolnowsky, a thread is selected from the single execution queue based on having the highest priority.

Browning teaches using a global execution queue to list threads that are executable. A dispatcher routine is used to select a thread from the global execution queue (col. 3, lines 27-34). Therefore, Browning fails to teach or suggest the limitations of ordering all requests from all of the hardware devices in a single queue and that all of the threads are simultaneously rearranged in the single queue, as recited in claims 1 and 9.

Jones teaches selecting threads for execution based on deadlines and is not relied on by the Examiner for disclosing the limitations of ordering all requests from all of the hardware devices in a single queue and that all of the threads are simultaneously rearranged in the single queue.

Ramakrishnan also fails to teach or suggest the limitations of ordering all requests from all of the hardware devices in a single queue and that all of the threads are simultaneously rearranged in the single queue. More specifically, Ramakrishnan teaches using a round-robin scheduler to select each thread for execution from a real time domain and a general purpose domain. Even if a single queue were used to store the threads for processing, as suggested by the Examiner, the order in which the threads would be serviced is determined by the round-robin scheduler using real time thread flags to select threads for execution (see Figure 2). Such an approach is different than that recited in the claims.

Zolnowsky also fails to teach or suggest the limitations of masking and unmasking interrupts from hardware devices. Zolnowsky teaches that each queue has a lock that must be acquired by a scheduler in order to dispatch a thread from the queue. With regard to interrupts, Zolnowsky teaches that "the holder of the schedule lock runs at an elevated interrupt level" (see col. 1, line 67 – col. 2, line 2). In contrast with the claimed mechanism of masking and unmasking interrupts from hardware devices recited in claims 1 and 9, the mechanism taught by Zolnowsky changes the interrupt level.

Regarding claims 2 and 14:

Claims 2 and 14 recite the limitations that latency information (or a hardware constraint) is computed based on a buffer size or display rate. The Examiner argues that col. 11, lines 24-47 of Ramakrishnan teaches this limitation. Ramadrishnan is completely silent regarding a display rate. Regarding a buffer size, Ramakrishnan simply states that it is "assumed that there is enough buffering...to avoid excessive dropping of packets because the buffers have filled." A maximum latency for processing a packet is determined based on a number of buffers. Nowhere does Ramakrishnan teach or suggest computing based on a buffer size, as expressly recited in claims 2 and 14.

Regarding claims 12 and 13:

Claim 12 recites the limitation that a hardware constraint is a size of a buffer and claim 13 recites the limitation that a hardware constraint is a fullness of a buffer. Further, as recited in claims 9 and 11, from which claims 12 and 13 both depend, the latency information used to schedule a thread is dependent on these hardware constraints. Therefore, the latency information used to schedule a thread is based on the size of the buffer or the fullness of the buffer. The Examiner argues that col. 11, lines 24-47 of Ramakrishnan teaches this limitation. However, as previously explained, the technique disclosed in Ramakrishnan assumes that there is enough buffering so there is no teaching or suggestion that the fullness or size of the buffering is a hardware constraint. The worst-case packet processing latency disclosed by Ramakrishnan is not computed based on either the size or fullness of a buffer, as recited in claims 12 and 13, respectively.

CONCLUSION

As the foregoing illustrates, no combination of the references cited by the Examiner can render independent claims 1 and 9 obvious. Since none of the references teach or suggest the limitations recited in claims 1 and 9, no combination of these references can render either claim 1 or claim 9 obvious. Since claims 2-8, 17-18, and 21-23 depend from allowable claim 1 and claims 10-15 depend from allowable claim 9, no combination of these references can render these claims obvious. Dependent claims 2 and 12-14 are also separately patentable for the more specific reasons set forth above with respect to these claims.

The Examiner errs in finding that:

- Claims 1, 3-4, 9-10, and 23 are unpatentable over Zolnowsky, in view of Browning, and further in view of Jones.
- Claims 1-15, 17-18, and 23 are unpatentable over Ramakrishnan, in view of Jones, and knowledge commonly known in the art, *as evidenced by* Browning.
- Claims 21 and 22 are unpatentable over Ramakrishnan, in view of Jones and Browning, and further in view of Cheng.

Respectfully submitted,

Stephanie Winner
Registration No. 52,371
PATTERSON & SHERIDAN, L.L.P.
3040 Post Oak Blvd. Suite 1500
Houston, TX 77056
Telephone: (650) 330-2310
Facsimile: (650) 330-2314
Agent for Applicant